

NVIDIA RAPIDS Python Workshop at NERSC for Future Perlmutter Users:

Welcome Talk

Rollin Thomas
DAS, NERSC
Tuesday 2020-04-13



Before We Begin

Thank you for Zooming in.

Please be sure to stay muted.

This is being recorded, should show up on NERSC website soon.


Slides are on the NERSC website.

You can use the Jupyter notebooks, GPU access through April 20.

For questions/comments, please use the chat window.

Stay relaxed, remain socially distanced, and wash your hands.

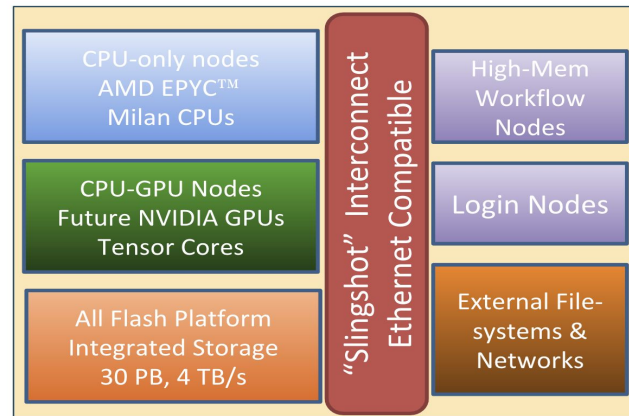
Now We Begin



Perlmutter is coming soon, > 3x computational power of Cori.
Heterogeneous CPU-node + GPU-node architecture.
New Cray Slingshot interconnect for data-centric computing.
All-flash platform integrated storage.

GPU nodes:

4 NVIDIA GPUs each with Tensor Cores, NVLink-3, HBM
1 AMD CPU
Unified Virtual Memory



GPUs, Python, and You



What do you need to know to make use of the GPUs?

Training events, workshops, hackathons.

This one is about data analytics on GPUs in Python with RAPIDS.

Why Python? Well,

>1000 unique non-staff Python users at NERSC

$\frac{1}{4}$ of all projects that run jobs at NERSC use Python

$\frac{1}{4}$ of all jobs that run at NERSC use Python somehow

Many EOD use cases leverage Python for productivity

Want to help *you* make use of Perlmutter GPUs on day one.

Who is Doing This



+



NVIDIA



**NERSC + NVIDIA/RAPIDS Partnership to
Prepare NERSC's Python Data Analytics Users for Perlmutter**

Format + Agenda

9:00-9:15	Welcome (right now)
9:15-9:30	Introduction to GPU Computing
9:30-10:15	Intro to RAPIDS, focused on cuDF and cuML
10:15-11:00	Introduction to cuDF <i>flipped classroom, notebooks sent in advance</i>
11:00-11:15	Break
11:15-12:00	Introduction to cuML <i>flipped classroom, notebooks sent in advance</i>
12:00-1:00	Break
1:00-1:30	Introduction to Dask <i>flipped classroom, notebooks sent in advance</i>
1:30-2:00	Introduction to Dask + GPUs
2:00-2:15	Evaluating CPU Workflows for the GPU <i>Thinking columnar, rather than row-wise</i>
2:15-3:15	Demo: Accelerate a Real Workflow
3:15-3:30	Break
3:30-4:00	Open Q&A, closing remarks, attendee survey

(Did you do the homework?)

Presenters from NVIDIA

Today's Presenters from NVIDIA:

Ayush Dattagupta

Vibhu Jawa

Zahra Ronaghi

Nick Becker

RAPIDS Engineering

RAPIDS Engineering

NVIDIA Solutions Architecture

RAPIDS Engineering

Presenters from NVIDIA

Today's Presenters from NVIDIA:

Ayush Dattagupta

Vibhu Jawa

Zahra Ronaghi

Nick Becker

RAPIDS Engineering

RAPIDS Engineering

NVIDIA Solutions Architecture

RAPIDS Engineering



THANK YOU SO MUCH



Further Thanks

Thanks also:

Laurie Stephey

NERSC

Seleste Rodriguez

NERSC

NERSC Users Like You

Socially Distanced



THANK YOU SO MUCH



Important links

<https://www.nersc.gov/users/training/events/rapids-hackathon/>

<https://github.com/beckernick/nersc-rapids-workshop>

<https://rapids.ai/>

<https://docs-dev.nersc.gov/cgpu/software/#python>

<https://docs.nersc.gov/programming/high-level-environments/python/>